Techniques for Developing Efficient Petascale Applications

Laxmikant (Sanjay) Kale

Department of Computer Science, University of Illinois at Urbana-Champaign

Urbana, Illinois, USA

PetaFLOPs class machine will start appearing during the coming year, and many multi-PetaFLOPs machines are on the anvil, including the NSF-funded sustained-PetaFLOPS track-1 machine, called *Blue Waters*, at NCSA/UIUC that will be operational in early 2011. It will be a substantial challenge to make the existing parallel CSE application codes run efficiently on them, and even more challenging to design new applications that can effectively leverage the large computational power of these machines.

In this talk I will first review the features of existing large supercomputers, including BlueGene/P, BlueGene/L and Cray XT4. I will then present a brief overview of the Blue Waters system, from the point of view of application developers.

Many techniques that are being developed (and that have been known from before) will be necessary to exploit these machines most effectively. I will present an overview of these, and illustrate them with application experience on existing machines with tens of thousands of processors.

Finally, I will describe *BigSim*, a system that allows early application development for future computers. In addition to helping identify scaling issues of data structures, it also allows application developers to identify potential performance bottlenecks of their applications on realistic data sets.